

Integrating Trust in Virtual Organisations^{*}

Ramón Hermoso, Holger Billhardt, and Sascha Ossowski

Artificial Intelligence Group

DATCCCIA - ESCET

University Rey Juan Carlos

{ramon.hermoso, holger.billhardt, sascha.ossowski}@urjc.es

Abstract. Organisational models cannot only be used to structure multiagent systems but also to express behaviour constraints for agents in open environments. However, sometimes these behaviour constraints cannot be exhaustively enforced, and some agents may transgress the norms put forward by a Virtual Organisation. This poses an additional burden on agents, as they cannot be sure that their acquaintances will behave as prescribed. Trust and reputation mechanisms are of particular relevance to this respect, as they are commonly used to infer expectations of future behaviour from past interactions.

In this paper we argue that, on the one hand, the a priori structure of Virtual Organisations can be useful to improve the efficiency of trust and reputation mechanisms, and that, on the other hand, such mechanisms provide relevant information for agents that are part of Virtual Organisations. For this purpose, we identify relevant aspects of existing organisational (meta-)models, and outline a reputation mechanism for Virtual Organisations that integrates these aspects. The dynamics of this mechanism is illustrated by an example.

1 Introduction

It is commonly agreed that the notion of organisation is of foremost importance to Multiagent Systems (MAS). In particular, organisational concepts are heavily used in the field of Agent-oriented Software Engineering [22]. In fact, it is tempting to maintain a tight coupling between a MAS, and the relevant features of the (human) organisation that it models, during the whole design process. Organisational concepts are often used as first-class abstractions that provide *structure* to the different models and stages of MAS design, and thus help designers to cope with high levels of complexity that MAS applications usually need to cope with [12].

When shifting the attention to *open* MAS, the *coercive* facets of organisational models gain relevance. Organisational abstractions are conceived as something aimed at limiting the freedom of choice of otherwise autonomous agents: once an agent freely chooses to enter an organisation in a certain position, playing certain roles, etc., it is supposed to behave in accordance with prescriptions

^{*} The present work was partly funded by the Spanish Ministry of Education and Science under grant TIC2003-08763-C02-02

attached to those concepts. Often, these prescriptions are complemented by a more general set of *norms* [19]. We refer to open MAS with these characteristics as *Virtual Organisations* (VOs) [16].

VOs differ in the way that prescriptions are enforced. Several approaches provide mechanisms to make it *impossible* for agents to transgress norms (e.g. by providing specific “governor” agents [4], or by integrating “filtering” mechanisms into MAS infrastructures [11]). However, especially for large-scale VOs this is a rather difficult and computationally expensive task. An alternative approach is to endow the VO with *incentive* mechanisms that, in general, make it too costly for agents to deviate from the prescribed behaviour (e.g. by means of installing incomplete but sufficiently effective detection and penalisation mechanisms for potential transgressors). Still, in the latter case, from the standpoint of an individual agent there is a significantly higher degree of uncertainty as to whether its organisational acquaintances will effectively behave in accordance with the organisational norms.

Several authors have investigated trust and reputation mechanisms that provide agents with *expectations* about the future behaviour of their acquaintances based on their interaction history within the MAS [8, 2, 15]. However, most mechanisms aim at supporting the emergence of overlay networks of trust relations in otherwise poorly structured systems. We believe that, on the one hand, the a priori structure of VOs can be useful to improve the efficiency of reputation mechanisms, and that, on the other hand, such mechanisms can be quite useful for an agent’s decision-making. This is particularly true for agents that are part of (and have to act to attain their individual goals within) VOs with “soft” enforcement mechanisms.

In this paper we present first results of our work in progress, aimed at integrating the structuring and coercive facets of organisational abstractions into trust and reputation mechanisms for VOs. In Section 2 we point to relevant aspects of previous work in the fields of organisational models and trust and reputation mechanisms. Section 3 outlines our proposal for building up and maintaining a trust model that takes into account organisational concepts, and shows how it can guide an agent’s decision-making in a VO. We present an example of how to apply those mechanisms within a particular VO in Section 4. Finally, we conclude summarising our proposal, compare it to approaches by other authors, and outline future lines of work.

2 Background

There is a wide range of organisational (meta-) models aimed at describing basic organisational concepts and their interrelation in the context of MAS [5, 7, 17, 10]. There is a common agreement that the notion of *role* is central for linking agents to an organisational model. Roles are sometimes defined by the *actions* they can perform, but usually they are characterised by the types of *social interactions* to which they contribute. The latter term does not primarily refer to the interaction protocols that agents engage in, but rather to the social

functionality that such interactions shall achieve. In this sense, we assume that VOs define roles and specify the interactions (functionalities) in which each role can participate.

Several meta-models allow for *specialisation* relations among essential organisational concepts. In the organisational model underlying the FIPA-ACL, for instance, *information exchange* interactions are a special kind of *request* interaction, where the requested action is a communicative action of type *inform* (e.g., [18]). In much the same way, the *informer* role involved in this interaction can be conceived as a specialisation of the *requester* role. In summary, organisational models often contain *taxonomies* of concept types, e.g., for roles or for interactions. Such taxonomies can be provided to the agents participating in an organisation – for instance, as an organisational service.

Finally, it is worth noting that many organisational models allow for certain types of *aggregation* relations. Different notions of *groups* – conceived as collections of agents – [5, 10], or collections of interaction protocols [6] are examples of such composed concepts. In a similar sense, the proper *organisation* itself as the aggregation of all its participating agents can be conceived as an individual unit.

There are many recent proposals for reputation mechanisms and approaches to evaluate trust in *peer-to-peer* systems in general (e.g. [21, 2]), and MAS in particular (e.g. [8, 20, 15]). Sabater and Sierra [14] consider reputation to have two different dimensions of influence: an *individual dimension* measuring local reputation – evaluated from direct interactions– and a *social dimension* evaluated from direct interactions and from the opinions from the society. In this paper, we will follow the proposal by Ramchurn et al. [13] regarding basic concepts of trust-based systems: *confidence* is a local rating based on direct interactions; *reputation* is a rating based on opinions of others; and *trust* is a rating built as a result from combining.

3 Trust Mechanisms in Virtual Organisations

Although VOs may limit the freedom of choice of agents, especially in less regulated organisations agents will still be confronted with the problem of deciding appropriate counterparts for their interactions according to their own beliefs and goals. Hence, trust and reputation mechanisms should be added as an additional layer on top of the organisational layer of a MAS as it is presented in Figure 1.

Not only a trust layer is useful for VOs. Also VOs provide a new viewpoint to trust and reputation mechanisms; organisational structures can help to get more reliable trust evaluations. In the following sections we present, first, an adaptation of standard trust and reputation mechanisms to VOs. Then we show how an agent can use knowledge about the organisational structure to infer confidence in an issue if no previous experience is available.




Fig. 1. Layered-network model for MAS

3.1 Basic Trust Model for Virtual Organisation

As described in Section 2, it seems reasonable to assume that a minimal organisational model defines at least roles and interactions, and that every agent participating in the VO plays at least one defined role. Furthermore, we assume that agents participating in an VO know the organisational structure, i.e. they know the roles other agents are playing within the organisation as well as the interactions that are defined for each role.

In line with other approaches [9, 20, 14, 13], we base our trust model on the notions of confidence and reputation. A typical situation is that an agent A wants to evaluate the trustworthiness of some other agent B – playing the role R – in the interaction I . This trustworthiness is denoted as $t_{A \rightarrow \langle B, R, I \rangle}$, with $t_{A \rightarrow \langle B, R, I \rangle} \in [0..1]$, and it measures the trust of A in B (playing role R) being a “good” counterpart in the interaction I . In order to build trust, agents can rely on two different measures: their own confidence, and the social reputation of an issue.

Confidence, $c_{A \rightarrow \langle B, R, I \rangle}$, is obtained from A ’s own experience when interacting with agent B playing role R in past interactions of type I . Confidence values for past interactions are stored in the agent’s *local interaction table* (LIT). This table contains one entry for each counterpart agent, playing a particular role, with which the agent has interacted in a particular interaction¹. LIT_A denotes agent A ’s LIT. An example is given in table 1.

Each entry in a LIT contains the following elements: i) the Agent/Role/Interaction identifier $\langle X, Y, Z \rangle$, ii) the confidence value for the issue ($c_{A \rightarrow \langle X, Y, Z \rangle}$), and iii) a reliability value ($r_{A \rightarrow \langle X, Y, Z \rangle}$). The confidence value may be obtained from some function that evaluates past experiences on the same issue. We suppose $c_{A \rightarrow \langle X, Y, Z \rangle} \in [0..1]$ and higher values to represent higher confidence. Reliability ($r_{A \rightarrow \langle X, Y, Z \rangle}$) measures how certain an agent is about its own

¹ Depending on computational restrictions, the table may resume all past events or just the recent interactions the agent was involved in.

$\langle X, Y, Z \rangle$	$c_{A \rightarrow \langle X, Y, Z \rangle}$	$r_{A \rightarrow \langle X, Y, Z \rangle}$
$\langle a_2, r_5, i_1 \rangle$	0.5	0.3
$\langle a_4, r_1, i_2 \rangle$	0.7	0.8
$\langle a_2, r_3, i_1 \rangle$	0.9	0.5
\vdots	\vdots	\vdots
$\langle a_9, r_2, i_5 \rangle$	0.4	0.7

Table 1. An agent's local interactions table (LIT_A)

confidence in issue $\langle X, Y, Z \rangle$. We suppose $r_{A \rightarrow \langle X, Y, Z \rangle} \in [0..1]$. Furthermore, we assume that $r_{A \rightarrow \langle X, Y, Z \rangle} = 0$ for any tuple $\langle B, R, I \rangle$ not belonging to LIT_A .

Reliability can be computed, for example, as proposed by Huynh, Jennings and Shadbolt [8, 9], by taking into account the number of interactions a confidence value is based on and the variability of the individual values across past experiences.

An agent may build trust directly from its confidence value or it may combine confidence with the social reputation of an issue. The latter is especially necessary if an agent has no experience on an issue or if its confidence is not sufficiently reliable. An agent can obtain the social reputation of an issue by asking other agents about their opinion on that issue. Agents that have been asked for their opinion return the corresponding confidence and reliability values from their LIT. Based on confidence and reputation, the trust A has in the issue $\langle B, R, I \rangle$ can be defined in the following way:

$$t_{A \rightarrow \langle B, R, I \rangle} = \begin{cases} c_{A \rightarrow \langle B, R, I \rangle}, & \text{if } r_{A \rightarrow \langle B, R, I \rangle} > \theta \\ \frac{\sum_{X \in RA} c_{X \rightarrow \langle B, R, I \rangle} \cdot w_{X \rightarrow \langle B, R, I \rangle}}{\sum_{X \in RA} w_{X \rightarrow \langle B, R, I \rangle}} & \text{otherwise} \end{cases} \quad (1)$$

Using this formula, trust will be measured at a scale $[0..1]$. θ is a threshold for the reliability of A 's own confidence values. If the reliability is below θ , $t_{A \rightarrow \langle B, R, I \rangle}$ is calculated as the weighted mean of the confidence values received from a set of *recommender agents* (RA). Agent A itself belongs to RA . $w_{X \rightarrow \langle B, R, I \rangle}$ is the weight given to agent X 's confidence on issue $\langle B, R, I \rangle$. This weight can be calculated as follows:

$$w_{X \rightarrow \langle B, R, I \rangle} = \begin{cases} r_{X \rightarrow \langle B, R, I \rangle} \cdot \alpha, & \text{if } X = A \\ r_{X \rightarrow \langle B, R, I \rangle} \cdot (1 - \alpha), & \text{otherwise} \end{cases} \quad (2)$$

where $\alpha \in [0..1]$ is a parameter specifying the importance given to A 's own confidence value. For values of $\alpha > 0.5$, an agent relies stronger on its own experience than on the opinions obtained from others.

One problem of reputation mechanisms is to determine the recommender agents that should be asked for their opinion about an issue $\langle B, R, I \rangle$. In a VO an agent A can take advantage of the organisational structures in order to decide which agents it should ask for their opinion. In fact, good recommenders may be

other agents that play the same role as A . The reason is twofold. First, agents playing the same role in an organisation will have the similar goals and, hence, it is likely that they have a similar subjective opinion about the trustworthiness of possible counterparts. Second, A probably wants to evaluate the trustworthiness of an issue $\langle B, R, I \rangle$ because A 's own role can participate in the interaction I . Thus, it is likely that other agents, playing the same role as A , will already have some experience with interaction I and possibly with the particular agent B .

3.2 Confidence and Trust for Organisational Structures

In the following sections we propose alternative ways to build an agent's confidence in an issue. We only concentrate on confidence values obtained from an agent's own experiences. The use of the proposed approaches in combination with social reputation in order to build trust is straight forward.

Agents accumulate past experiences in form of atomic confidence values for Agent/Role/Interaction tuples in their LIT. This information may be used to calculate confidence (and trust) values for other organisational elements by accumulating the corresponding entries in an agent's LIT.

Agent/role confidence (denoted by $c_{A \rightarrow \langle B, R, _ \rangle}$) evaluates an agent's trust in an other agent playing a specific role within the organisation and regardless of any concrete interaction. It measures the confidence an agent A has in agent B playing a role R and is calculated by aggregating all past experiences from agent A in any interaction where A met B (playing role R):

$$c_{A \rightarrow \langle B, R, _ \rangle} = \frac{\sum_{\langle X, Y, Z \rangle \in E} c_{A \rightarrow \langle X, Y, Z \rangle} \cdot r_{A \rightarrow \langle X, Y, Z \rangle}}{\sum_{\langle X, Y, Z \rangle \in E} r_{A \rightarrow \langle X, Y, Z \rangle}}, \quad (3)$$

where $E = \{\langle X, Y, Z \rangle \in LIT_A \mid X = B, Y = R\}$ is the set of all entries in A 's LIT with information about agent B playing role R (possibly within different interactions). In this way, $c_{A \rightarrow \langle B, R, _ \rangle}$ is calculated as the weighted mean of the confidence values for all issues in A 's LIT that refer to the agent B and the role R . The used weights are the corresponding reliability values.

Agent/role confidence may be used as an additional evidence measure when calculating $t_{A \rightarrow \langle B, R, I \rangle}$. However, more importantly it provides a manner to evaluate $c_{A \rightarrow \langle B, R, I \rangle}$ (and $t_{A \rightarrow \langle B, R, I \rangle}$) if agent A has none or not enough experience regarding the issue $\langle B, R, I \rangle$, that is, if $r_{A \rightarrow \langle B, R, I \rangle} < \theta$. The importance increases if none of the agents in the organisation has had any experience regarding the issue $\langle B, R, I \rangle$, and therefore, none of the agents could give any (reliable) recommendation. In such a scenario, $c_{A \rightarrow \langle B, R, _ \rangle}$ can provide a valuable approximation of $c_{A \rightarrow \langle B, R, I \rangle}$ for any interaction I .

In a similar way, agents can compute *agent confidence* ($c_{A \rightarrow \langle B, _ \rangle}$) – the (global) confidence agent A has in agent B . Agent confidence values can provide a second level of approximation when building $t_{A \rightarrow \langle B, R, I \rangle}$. They may be used as an alternative for $c_{A \rightarrow \langle B, R, I \rangle}$ if there is not even enough expertise for a reliable

confidence $c_{A \rightarrow \langle B, R, _ \rangle}$. In a more general environment with agents possibly participating in several organisations, agent confidence may also be used as a gauge to authorise agents to join an organisation.

Equation 3 can be adapted to calculate *role confidence* ($c_{A \rightarrow \langle _, R, _ \rangle}$) and *interaction confidence* ($c_{A \rightarrow \langle _, _, I \rangle}$)². Role confidence measures an agents confidence in a specific role within an organisation. It could be used as a default confidence value assigned to agents that just entered an organisation playing a specific role and, thus, for which there are no confidence values available. Interaction confidence provides an estimation of the trust in a concrete interaction within an organisation despite the actual agents that have participated in the interaction. Although interactions are defined from specific sets of roles, interaction confidence can be calculated independently of roles by aggregating all entries related to a specific interaction in an agent's LIT. Interaction confidence may be used as a means to choose between several alternative interactions an agent could participate in.

Role and interaction confidence have an additional importance for VOs. They evaluate certain parts of an organisation at the institutional level, that is, independently on the agents actually participating in the organisation. From the institutional point of view, role and interaction confidence can be used to identify deficiencies in the organisational structure and functioning. From the outside, both measures can be used to evaluate parts of an organisation or an organisation as a whole.

Confidence (and trust) values can also be aggregated for groups of agents – either in general or in relation to one or more interactions or roles. Suppose, an agent A intends to evaluate a group of agents $AG = \{B_1, B_2, \dots, B_n\}$, for instance, in order to decide whether to join them or not. The value can be estimated from A 's experience about past interactions in which also agents belonging to the group AG participated. The confidence an agent A has in the group AG with regard to an interaction I can be obtained by aggregating all experiences of A about the behaviour of agents belonging to AG in the interaction type I . Formally, the value can be calculated as the mean of all of A 's confidence values for tuples $\langle X, Y, Z \rangle$ with $X \in AG$ and $Z = I$. This mean is weighted by the corresponding reliability values. The formula is as follows:

$$c_{A \rightarrow \langle AG, _, I \rangle} = \frac{\sum_{\langle X, Y, Z \rangle \in E} c_{A \rightarrow \langle X, Y, Z \rangle} \cdot r_{A \rightarrow \langle X, Y, Z \rangle}}{\sum_{\langle X, Y, Z \rangle \in E} r_{A \rightarrow \langle X, Y, Z \rangle}}, \quad (4)$$

where $E = \{\langle X, Y, Z \rangle \in LIT_A \mid Z = I \wedge X \in AG\}$ is the set of all entries in A 's LIT with information about agents belonging to group AG that participate in interaction I , possibly playing different roles.

In a similar way it is possible to compute the confidence in a group of agents in general ($c_{A \rightarrow \langle AG, _, _ \rangle}$), or in a group of agents with respect to a specific set of roles and/or interactions.

² It is also possible to compute *agent/interaction confidence* values ($c_{A \rightarrow \langle B, _, I \rangle}$). However, we do not consider this measure very useful.

Finally, considering AG to be the set of all agents participating in a specific organisation, then $c_{A \rightarrow \langle AG, \dots \rangle}$ specifies the confidence in that organisation – measured as the aggregation of the confidence values for all participating agents. Organisation confidence may have a special importance in open MAS with where several VOs compete with each other.

3.3 Confidence Inference using Role and Interaction Similarities

In the previous subsection we have proposed to use the agent/role confidence $c_{A \rightarrow \langle B, R, \dots \rangle}$ (or $c_{A \rightarrow \langle B, \dots, \dots \rangle}$) as an estimation for $c_{A \rightarrow \langle B, R, I \rangle}$ if agent A has no reliable experience about issue $\langle B, R, \dots \rangle$. This approach is based on the hypothesis that, in general, agents behave in a similar way in all interactions related to the same role. Formally, we assume that for any interaction I' , with $I' \neq I$, the value $c_{A \rightarrow \langle B, R, I' \rangle}$ is an approximation for $c_{A \rightarrow \langle B, R, I \rangle}$. Refining this idea, it seems reasonable to assume that the more similar I' and I the more similar will be the values $c_{A \rightarrow \langle B, R, I' \rangle}$ and $c_{A \rightarrow \langle B, R, I \rangle}$. And the same actually applies to roles. Using this idea, confidence values accumulated for similar agent/role/interaction tuples may provide evidence for the value of $c_{A \rightarrow \langle B, R, I \rangle}$. We propose the following equation for calculating confidence:

$$c_{A \rightarrow \langle B, R, I \rangle} = \frac{\sum_{\langle X, Y, Z \rangle \in E} c_{A \rightarrow \langle X, Y, Z \rangle} \cdot r_{A \rightarrow \langle X, Y, Z \rangle} \cdot \text{sim}(\langle X, Y, Z \rangle, \langle B, R, I \rangle)}{\sum_{\langle X, Y, Z \rangle \in E} r_{A \rightarrow \langle X, Y, Z \rangle} \cdot \text{sim}(\langle X, Y, Z \rangle, \langle B, R, I \rangle)} \quad (5)$$

Using equation 5, each entry from agent A 's LIT has an influence in the calculation of $c_{A \rightarrow \langle B, R, I \rangle}$. The weight given to an entry is determined by the similarity of the agent/role/interaction key to the key $\langle B, R, I \rangle$ and by the reliability of the confidence value. $\text{sim}(\langle X, Y, Z \rangle, \langle B, R, I \rangle)$ can be computed as the weighted sum of the similarities of the individual elements (agent, role, and interaction), as defined in the following equation:

$$\text{sim}(\langle X, Y, Z \rangle, \langle B, R, I \rangle) = \alpha \cdot \text{sim}_A(B, X) + \beta \cdot \text{sim}_R(R, Y) + \gamma \cdot \text{sim}_I(I, Z) \quad (6)$$

where $\text{sim}_A(B, X)$, $\text{sim}_R(R, Y)$, $\text{sim}_I(I, Z) \in [0..1]$ measure the similarity between the agents, roles and interactions, respectively. α, β and γ , with $\alpha + \beta + \gamma = 1$, are parameters specifying the sensibility regarding the individual similarities. Assuming that only confidence values for the same agent are taken into account, $\text{sim}_A(B, X)$ is defined as follows:

$$\text{sim}_A(B, X) = \begin{cases} 1, & \text{if } B = X \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

As argued in Section 2, many organisational models include taxonomies of roles and/or interactions. If this is the case, $\text{sim}_R(R, R')$ and $\text{sim}_I(I, I')$ can be implemented by *closeness functions* that estimate the similarity between two concepts on the basis of their closeness in the concept hierarchy.

Equations 3 and 5 can be used as an additional indicator for $t_{A \rightarrow \langle B, R, I \rangle}$. If an agent has no reliable experience about a particular agent/role/interaction

issue, they can be used to estimate trust without the necessity to rely on the opinions of other agents. Thus, the proposed model makes agents less dependent on others, which is an important issue especially in VOs that do not provide mechanisms to keep its members from cheating.

4 An Example

In this section we illustrate our approach with an example taken from the University domain. Suppose a School of Computer Science whose members play roles out of the taxonomy shown in Figure 2. Furthermore, suppose that the social functionalities provided by the organisation are summarised in the interaction taxonomy illustrated in Figure 3.




Fig. 2. Role taxonomy provided by University organisation

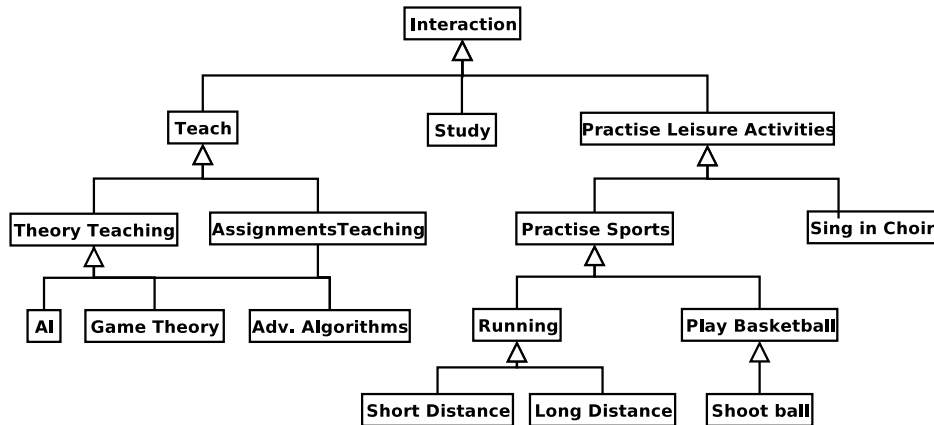


Fig. 3. Interaction taxonomy provided by University organisation

Suppose an agent a , playing the role *Student*, has just finished its first year at the university and wants to enrol in the second year. Suppose a wants to enrol

in the subject *Artificial Intelligence (AI)* and suppose there are two different lecturers giving AI classes: L_1 and L_2 . Agent a has to decide one of those lecturers to attend his/her classes. In order to make this decision agent a can use its experience stored in form of confidence values in its LIT. This LIT is shown in Table 2.

$\langle X, Y, Z \rangle$	$c_{a \rightarrow \langle X, Y, Z \rangle}$	$r_{a \rightarrow \langle X, Y, Z \rangle}$
$\langle L_7, Th.Lecturer, Teach \rangle$	0.1	0.3
$\langle L_2, Th.Lecturer, Teach_GameTheory \rangle$	0.9	0.3
$\langle S_7, Student, Study \rangle$	0.7	0.5
$\langle S_7, Sprinter, Running \rangle$	0.3	0.7
$\langle L_1, Ass.Lecturer, Teach_AdvancedAlgorithms \rangle$	0.8	0.9
$\langle L_2, Lecturer, Teach_GameTheoryAssignments \rangle$	0.3	0.8
$\langle L_1, Th.Lecturer, Teach \rangle$	0.7	0.5
$\langle L_1, Singer, Sing_In_Choir \rangle$	0.5	0.5
$\langle S_4, Student, Study \rangle$	0.7	0.8

Table 2. Agent a 's local interaction table

As it can be seen in Table 2, agent a has no direct experience, that is, no confidence values, for the issues $\langle L_1, Lecturer, Teach_{AI} \rangle$ and $\langle L_2, Lecturer, Teach_{AI} \rangle$. However, using knowledge about the organisation, a could calculate $c_{a \rightarrow \langle L_1/L_2, Lecturer, _ \rangle}$ as an approximation for the desired confidence values. In doing so, the agent obtains:

$$c_{a \rightarrow \langle L_1, Lecturer, _ \rangle} = novalue \quad c_{a \rightarrow \langle L_2, Lecturer, _ \rangle} = 0.3.$$

Observe that there is no possible rating for $c_{a \rightarrow \langle L_1, Lecturer, _ \rangle}$ since there is no matching between tuple $\langle L_1, Lecturer, _ \rangle$ and any other tuple in LIT_A .

In this case, as a second approximation, agent a could calculate agent confidence:

$$c_{a \rightarrow \langle L_1, _ \rangle} = 0.56 \quad c_{a \rightarrow \langle L_2, _ \rangle} = 0.46.$$

This approximation, however, would not only count the confidence in L_1 and L_2 as teachers, but also, for example, the confidence in L_1 as a *Singer* in the *Choir*. Hence, it would be better to calculate $c_{a \rightarrow \langle L_1/L_2, Lecturer, Teach_{AI} \rangle}$ using equation 5. In order to do that, the agent could use the following simple equation to calculate the similarity between roles and interactions, respectively:

$$sim_R(x, y) = sim_I(x, y) = 1 - \frac{h}{h_{MAX}} \quad (8)$$

where x, y are either roles or interactions, h is the number of nodes between x and y in the taxonomy, and h_{MAX} is the longest possible path between any pair of elements in the hierarchy tree.

Using equations 5 and 8 the calculated confidence values are the following:

$$c_{a \rightarrow \langle L_1, Lecturer, Teach_{AI} \rangle} = 0.61 \quad c_{a \rightarrow \langle L_2, Lecturer, Teach_{AI} \rangle} = 0.54,$$

(We have set $\alpha = 0.5$, $\beta = 0.35$ and $\gamma = 0.15$.) Based on these values, agent a decides to attend the classes from lecturer L_1 .

5 Conclusion

In this paper we have presented results of our work in progress, aimed at integrating organisational facets into trust mechanisms. We have emphasised the problem of finding “good” counterparts, even if no previous interactions have been performed. The proposed trust model takes into account key concepts of organisational models, such as *roles* and *interactions*, as well as their aggregation in *groups* or *organisations*. We have also endowed our model with inference capabilities exploiting *taxonomies* of concept types provided by VOs.

In contrast to other approaches to trust systems (most of them based on reputation distribution), we have presented a way of evaluating trust at a local level that emphasizes the different experiences of agents from past interactions. The FIRE model proposed by Huynh, Jennings and Shadbald [8] is also concerned with *interaction trust* and *role-based trust*. As in our approach, the former is built from direct experience of an agent, while the latter is the rating that results from role-based relationships between agents. Nevertheless, the FIRE model does not consider inference on VO structures.

The model proposed by Sabater and Sierra [15] also exploits ontologies to make up trust values. Nevertheless, it does not consider organisations as a significant concept, and thus it does not take into account organisational structures.

Abdul-Rahman and Hailes [1] propose a trust model for virtual communities but use qualitative ratings for estimating trust. They focus on evaluating trust from past expertise and reputation coming from *recommender agents* without considering explicitly VO structures.

The trust model by Ramchurn et al. [13] is based on direct and indirect multi-agent interactions for establishing contracts between agents in electronic institutions[3]. Still, it does not account for VOs with “soft” enforcement mechanisms, where norms and behaviour rules can be transgressed.

We are currently finishing a testbed that will allow us to simulate VOs with the proposed trust model, so as to gain experimental evidence regarding its behaviour in different situations. It will also allow us to come up with a more quantitative comparison to other approaches. Finally, we will investigate as to how far existing organisational MAS methodologies can be extended so as to account for trust mechanisms in VOs.

References

1. Alfarez Abdul-Rahman and Stephen Hailes. Supporting trust in virtual communities. In *HICSS*, 2000.
2. Karl Aberer and Zoran Despotovic. Managing trust in a peer-2-peer information system. In *CIKM '01: Proceedings of the tenth international Conference on Information and Knowledge Management*, pages 310–317, New York, NY, USA, 2001. ACM Press.

3. M. Esteva, J. A. Rodríguez, C. Sierra, P. Garcia, and J. L. Arcos. On the formal specifications of electronic institutions. In F. Dignum and C. Sierra, editors, *Agent-mediated Electronic Commerce (The European AgentLink Perspective)*, volume 1191 of *LNAI*, pages 126–147, Berlin, 2001. Springer.
4. Marc Esteva, Bruno Rosell, Juan A. Rodríguez-Aguilar, and Josep Ll. Arcos. AMELI: An agent-based middleware for electronic institutions. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 236–243, 2004.
5. J. Ferber and O. Gutknecht. A meta-model for the analysis of organizations in multi-agent systems. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS'98)*, pages 128–135, Paris, France, July 1998. IEEE Press.
6. Ramón Hermoso, Rubén Ortiz, Sergio Saugar, and J. M. Serrano. Instrumentación de sistemas multi-agente mediante un entorno organizativo/lingüístico: Un caso práctico. In Juan Carlos González Moreno, Pedro Cuesta Morales, and Juan Pavón Mestras, editors, *I Taller en Desarrollo de Sistemas Multiagente (DESMA-2004)*, pages 72–83, Málaga, Nov 2004.
7. J. Fred Hübner, J. Simão Sichman, and O. Boissier. Using the moise+ for a cooperative framework of mas reorganisation. In *Advances in Artificial Intelligence - SBIA 2004*, pages 506–515, 2004.
8. T. Dong Huynh, Nicholas R. Jennings, and Nigel R. Shadbolt. Developing an integrated trust and reputation model for open multi-agent systems. In Rino Falcone, Suzanne Barber, Jordi Sabater, and Munindar Singh, editors, *AAMAS-04 Workshop on Trust in Agent Societies*, 2004.
9. T. Dong Huynh, Nicholas R. Jennings, and Nigel R. Shadbolt. FIRE: An integrated trust and reputation model for open multi-agent systems. In *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI)*, 2004.
10. James Odell, Marian H. Nodine, and Renato Levy. A metamodel for agents, roles, and groups. In James Odell, Paolo Giorgini, and Jörg P. Müller, editors, *AOSE*, volume 3382 of *Lecture Notes in Computer Science*, pages 78–92. Springer, 2004.
11. A. Omicini, S. Ossowski, and A. Ricci. Coordination infrastructures in the engineering of multiagent systems. In Federico Bergenti, Marie-Pierre Gleizes, and Franco Zambonelli, editors, *Methodologies and Software Engineering for Agent Systems: The Agent-Oriented Software Engineering Handbook*, chapter 14, pages 273–296. Kluwer Academic Publishers, June 2004.
12. Andrea Omicini and Sascha Ossowski. Objective versus subjective coordination in the engineering of agent systems. In Matthias Klusch, Sonia Bergamaschi, Peter Edwards, and Paolo Petta, editors, *Intelligent Information Agents: An AgentLink Perspective*, volume 2586 of *LNAI: State-of-the-Art Survey*, pages 179–202. Springer-Verlag, March 2003.
13. Sarvapali D. Ramchurn, Carles Sierra, Lluís Godó, and Nicholas R. Jennings. A computational trust model for multi-agent interactions based on confidence and reputation. In *Proceedings of 6th International Workshop of Deception, Fraud and Trust in Agent Societies*, pages 69–75, 2003.
14. Jordi Sabater and Carles Sierra. REGRET: a reputation model for gregarious societies. In Jörg P. Müller, Elisabeth Andre, Sandip Sen, and Claude Frasson, editors, *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 194–195, Montreal, Canada, 2001. ACM Press.
15. Jordi Sabater and Carles Sierra. Reputation and social network analysis in multi-agent systems. In *AAMAS '02: Proceedings of the first international joint confer-*

- ence on Autonomous agents and multiagent systems*, pages 475–482, New York, NY, USA, 2002. ACM Press.
16. M. Schumacher and S. Ossowski. The governing environment. In Weyns, Parunak, and Michel, editors, *Environments for Multiagent Systems II*, volume 3830, pages 88–104. Springer-Verlag, 2006.
 17. Juan M. Serrano, S. Ossowski, and A. Fernández. The pragmatics of software agents - analysis and design of agent communication languages. *Intelligent Information Agents - An AgentLink Perspective (Klusch, Bergamaschi, Edwards & Petta, ed.), Lecture Notes in Computer Science*, 2586:234–274, 2003.
 18. Juan M. Serrano and Sascha Ossowski. On the impact of agent communication languages on the implementation of agent systems. *Cooperative Information Agents VIII (Klusch, Ossowski & Kashyap, ed.)*, 2004.
 19. Javier Vázquez-Salceda, Virginia Dignum, and Frank Dignum. Organizing multiagent systems. *Autonomous Agents and Multi-Agent Systems*, 11(3):307–360, 2005.
 20. Bin Yu and Munindar P. Singh. An evidential model of distributed reputation management. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 294–301, New York, NY, USA, 2002. ACM Press.
 21. Bin Yu, Munindar P. Singh, and Katia Sycara. Developing trust in large-scale peer-to-peer systems. In *Proceedings of First IEEE Symposium on Multi-Agent Security and Survivability*, pages 1–10, 2004.
 22. Franco Zambonelli, Nicholas R. Jennings, and Michael Wooldridge. Organizational abstractions for the analysis and design of multi-agent systems. In Paolo Ciancarini and Michael J. Wooldridge, editors, *AOSE*, volume 1957 of *LNCS*, pages 235–252. Springer, 2000.